FFPACK: Finite Field Linear Algebra Package

Jean-Guillaume Dumas
Université Joseph Fourier,
Lab. de Modélisation et Calcul,
50 av. des Mathématiques.
B.P. 53 X, 38041 Grenoble,
France.
Jean-Guillaume.Dumas@imag.fr

Pascal Giorgi
Lab. de l'Informatique du Parallélisme,
ENS Lyon. 46, allée d'Italie,
F69364 Lyon Cédex 07,
France.
Pascal.Giorgi@ens-lyon.fr

Clément Pernet
Université Joseph Fourier,
Lab. de Modélisation et Calcul,
50 av. des Mathématiques.
B.P. 53 X, 38041 Grenoble,
France.
Clement.Pernet@imag.fr

ABSTRACT

The FFLAS project has established that exact matrix multiplication over finite fields can be performed at the speed of the highly optimized numerical BLAS routines. Since many algorithms have been reduced to use matrix multiplication in order to be able to prove an optimal theoretical complexity, this paper shows that those optimal complexity algorithms, such as LSP factorization, rank determinant and inverse computation can also be the most efficient.

Categories and Subject Descriptors

G.4 [**Mathematical Software**]: Algorithm design and analysis; F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Computations in finite fields.

General Terms

Algorithms, Experimentation, Performance.

Keywords

Finite fields; BLAS; LSP Factorization

1. INTRODUCTION

Exact matrix multiplication over finite fields can now be performed at the speed of the highly optimized numerical BLAS routines. This has been established by the FFLAS project [8]. Moreover, since finite field computations e.g. do not suffer from numerical stability, this project showed also an easy effectiveness of the algorithms with even better arithmetic complexity (such as Winograd's variant of Strassen's fast matrix multiplication) [18].
Now for the applications. Many algorithms have been designed to use matrix multiplication in order to be able to prove an optimal theoretical complexity. In practice those algorithms were only seldom used. This is the case e.g. in many linear algebra problems such as determinant, rank, inverse, system solution or minimal and characteristic polynomial. Over finite fields or over the integers those finite field linear algebra routines are used to solve many different problems. Among them are integer polynomial factorization, Gröbner basis computation, integer system solving, large integer factorization, discrete logarithms, error correcting codes, etc. Even sparse or polynomial linear algebra needs some very efficient dense subroutines [13, 11]. We believe that with our kernel, each one of those optimal complexity algorithms can also be the most efficient.
The goal of this paper is to show the actual effectiveness of this belief for the factorization of any shape and any rank matrices. The application of this factorization to determinant, rank, and inverse is presented as well. Some of the ideas from FFLAS, in particular the fast matrix multiplication algorithm for small prime fields, are now incorporated into the Maple computer algebra system since its version 8. Therefore an effort towards effective reduction has been made within Maple by A. Storjohann[4]. Effective reduction for minimal and characteristic polynomial were sketched in [20] and A. Steel has reported on similar efforts within his implementation of some Magma routines. We provide a full C++ package available directly[1] or through LinBox[2][7]. Extending the work undertaken by the authors et al.[18, 8, 3, 12], this paper focuses on matrix factorization, namely the exact equivalent of the LU factorization. Indeed, unlike numerical matrices, exact matrices are very often singular, even more so if the matrix is not square ! Consequently, Ibarra, Moran and Hui have developed generalizations of the LU factorization, namely the LSP and LQUP factorizations [16]. In section 4 we deal with the implementation of those two routines as well as memory optimized (an in-place and a cyclic block) versions using the fast matrix multiplication kernel. Those implementations require the resolution of triangular systems with matrix right or left hand side. In section 3 the triangular system solver (`Trsm` routine using BLAS terminology) is therefore studied. Then, in section 5, we propose different uses of the factorization routine to solve other classical linear algebra problems. In particular, speed ratios are presented and reflects the optimal behavior of our routines.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ISSAC'04, July 4–7, 2004, Santander, Spain.
Copyright 2004 ACM 1-58113-827-X/04/0007 ...$5.00.

[1] `www-lmc.imag.fr/lmc-mosaic/Jean-Guillaume.Dumas/FFLAS`
[2] `www.linalg.org`

2. BASE FIELDS

The algorithms we present in this paper are written generically with regards to the field over which they operate, as long as they provide some conversion functions from a field element to a floating point representation and backwards. As demonstrated in [8] this is easily done for prime fields and also for other finite fields, via a q-adic transformation. The chosen interface is that of the LinBox fields [21, §5.3]. In the following, the prime field with characteristic p will be denoted by $\mathbb{Z}_p$.

For our experiments we use some classical representations, e.g. modular prime fields, primitive roots Galois fields, Montgomery reduction, etc. implemented in different libraries, as in [8, 6]. Still and all, when no special implementation is required and when the prime field is small enough, one could rather use what we call a `Modular<double>` field representation. Indeed, the use of the BLAS imposes conversions between the field element representations and a corresponding floating point representation. Hence a lot of time consuming conversions can be avoided whenever the field element representation is already a floating point number. This is the case for the LinBox prime field `Modular<double>`, where the exact representation of an element is stored within the mantissa of a double precision floating point number. Of course all the arithmetic operations remain exact as they are always performed modulo a prime number. In this paper, we only focus on two representations. The first one is `Modular<double>` and the second one is that of the Givaro library [3] which uses machine integer remaindering. This classical representation will be denoted by `Givaro-ZpZ`.

3. TRIANGULAR SYSTEM SOLVING WITH MATRIX HAND SIDE

In this section we discuss the implementation of solvers for triangular systems with matrix right hand side (or equivalently left hand side). This is also the simultaneous resolution of n triangular systems. Without loss of generality for the triangularization, we here consider only the case where the row dimension, m, of the the triangular system is less than or equal to the column dimension, n. The resolution of such systems is a classical problem of linear algebra. It is e.g. one of the main operation in block Gaussian elimination. For solving triangular systems over finite fields, the block algorithm reduces to matrix multiplication and achieves the best known arithmetic complexity. Therefore, from now on we will denote by ω the exponent of square matrix multiplication (e.g. from 3 for classical, to 2.375477 for Coppersmith-Winograd). Moreover, we can bound the arithmetical cost of a $m \times k$ by $k \times n$ rectangular matrix multiplication (denoted by $R(m,k,n)$) as follows: $R(m,k,n) \le C_\omega min(m,k,n)^{\omega-2} max(mk,mn,kn)$ [15, (2.5)]. In the following subsections, we present the block recursive algorithm and two optimized implementation variants.

3.1 Scheme of the block recursive algorithm

The classical idea is to use the divide and conquer approach. Here, we consider the upper left triangular case without loss of generality, since the any combination of upper/lower and left/right triangular cases are similar: if U is upper triangular, L is lower triangular and B is rectangular, we call `ULeft-Trsm` the resolution of $UX = B$, `LLeft-Trsm` that of $LX = B$, `URight-Trsm` that of $XU = B$ and `LRight-Trsm` that of $XL = B$.

Algorithm `ULeft-Trsm`(A, B)
Input: $A \in \mathbb{Z}_p{}^{m\times m}$, $B \in \mathbb{Z}_p{}^{m \times n}$.
Output: $X \in \mathbb{Z}_p{}^{m\times n}$ such that $AX = B$.
Scheme
if m=1 **then**
$X := A_{1,1}^{-1} \times B$.
else *(splitting matrices into $\lfloor \frac{m}{2} \rfloor$ and $\lceil \frac{m}{2} \rceil$ blocks)*

$$\overbrace{\begin{bmatrix} A_1 & A_2 \\ & A_3 \end{bmatrix}}^{A} \overbrace{\begin{bmatrix} X_1 \\ X_2 \end{bmatrix}}^{X} = \overbrace{\begin{bmatrix} B_1 \\ B_2 \end{bmatrix}}^{B}$$

$X_2 :=$`ULeft-Trsm`(A_3, B_2).
$B_1 := B_1 - A_2X_2$.
$X_1 :=$`ULeft-Trsm`(A_1, B_1).
return X.

Lemma 3.1. *Algorithm* `ULeft-Trsm` *is correct and its theoretical cost is bounded by* $\frac{C_\omega}{2(2^{\omega-2}-1)} nm^{\omega-1}$ *arithmetic operations in* $\mathbb{Z}_p$ *for* $m \le n$.

Proof. The correctness of algorithm `ULeft-Trsm` can be proven by induction on the row dimension of the system. For this, one only has to note that

$$X = \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} \text{is solution} \iff \begin{cases} A_3X_2 = B_2 \\ A_1X_1 + A_2X_2 = B_1 \end{cases}$$

Let $C(m,n)$ be the cost of algorithm `ULeft-Trsm` where m is the dimension of A and n the column dimension of B. It follows from the algorithm that $C(m,n) = 2C(\frac{m}{2}, n) + R(\frac{m}{2}, \frac{m}{2}, n)$. By counting each operation at one recursive step we have:

$$C(m,n) = \sum_{i=1}^{\log m} 2^{i-1} R(\frac{m}{2^i}, \frac{m}{2^i}, n)$$

Now, since $m \le n$, we get $\forall i\ R(\frac{m}{2^i}, \frac{m}{2^i}, n) = C_\omega \left(\frac{m}{2^i}\right)^{\omega-1} n$ and therefore:

$$C(m,n) = \frac{C_\omega nm^{\omega-1}}{2} \sum_{i=1}^{\log m} \left(\frac{1}{2^i}\right)^{\omega-2}$$

which gives the $O(nm^{\omega-1})$ bound of the lemma. □

3.2 Implementation using the BLAS "dtrsm"

Matrix multiplication speed over finite fields was improved in [8, 18] by the use of the numerical BLAS[4] library: matrices were converted to floating point representations (where the linear algebra routines are fast) and converted back to a finite field representation afterwards. The computations remained exact as long as no overflow occurred. An implementation of `ULeft-Trsm` can use the same techniques. Indeed, as soon as no overflow occurs one can replace the recursive call to `ULeft-Trsm` by the numerical BLAS *dtrsm* routine. But one can remark that approximate divisions can occur. So we need to ensure both that only exact divisions are performed and that no overflow appears. Not only one

[3] `www-id.imag.fr/Logiciels/givaro`

[4] `www.netlib.org/blas`

has to be careful for the result to remain within acceptable bounds, but, unlike matrix multiplication where data grows linearly, data involved in linear system grows exponentially as shown in the following.
The next two subsections first show how to deal with divisions, and then give an optimal theoretical bound on the coefficient growth and therefore an optimal threshold for the switch to the numerical call.

3.2.1 Dealing with divisions

In algorithms like `ULeft-Trsm`, all divisions appear only within the last recursion's level. In the general case it cannot be predicted whether these divisions will be exact or not. However when the system is unitary (only 1's on the main diagonal) the division are of course exact and will even never be performed. Our idea is then to transform the initial system so that all the recursive calls to `ULeft-Trsm` are unitary. For a triangular system $AX = B$, it suffices to factor first the matrix A into $A = UD$, where U, D are respectively an upper unit triangular matrix and a diagonal matrix. Next the unitary system $UY = B$ is solved by any `ULeft-Trsm` (even a numerical one), without any division. The initial solution is then recovered over the finite field via $X = D^{-1}Y$. This normalization leads to an additional cost of:

- m inversions over $\mathbb{Z}_p$ for the computation of D^{-1}.
- $(m-1)\frac{m}{2} + mn$ multiplications over $\mathbb{Z}_p$ for the normalizations of U and X.

Nonetheless, in our case, we need to avoid divisions only during the numerical phase. cwTherefore, the normalization can take place only just before the numerical routine calls. Let β be the size of the system when we switch to a numerical computation. To compute the cost, we assume that $m = 2^i\beta$, where i is the number of recursive level of the algorithm `ULeft-Trsm`. The implementation can however handle any matrix size. Now, there are 2^i normalizations with systems of size β. This leads to an additional cost of:

- m inversions over $\mathbb{Z}_p$.
- $(\beta - 1)\frac{m}{2} + mn$ multiplications over $\mathbb{Z}_p$.

This allows us to save $\left(\frac{1}{2} - \frac{1}{2^{i+1}}\right) m^2$ multiplications over $\mathbb{Z}_p$ from a whole normalization of the initial system. One iteration suffices to save $\frac{1}{4}m^2$ multiplications and we can save up to $\frac{1}{2}(m^2 - m)$ multiplications with $\log m$ iterations.

3.2.2 A theoretical threshold

The use of the BLAS routine trsm is the resolution of the triangular system over the integers (stored as `double` for `dtrsm` or `float` for `strsm`). The restriction is the coefficient growth in the solution. Indeed, the k^{th} value in the solution vector is a linear combination of the $(n-k)$ already computed next values. This implies a linear growth in the coefficient size of the solution, with respect to the system dimension. Now this resolution can only be performed if every element of the solution can be stored in the mantissa of the floating point representation (e.g. 53 bits for `double`). Therefore overflow control consists in finding the largest block dimension b, such that the result of the call to `dtrsm` will remain exact.

We now propose a bound for the values of the solutions of such a system; this bound is optimal (in the sense that there exists a worst case matching the bound when $n = 2^i b$). This enables the implementation of a cascading algorithm, starting recursively and taking advantage of the BLAS performances as soon as possible.

Theorem 3.2. *Let $T \in \mathbb{Z}^{n\times n}$ be a unit diagonal upper triangular matrix, and $b \in \mathbb{Z}^n$, with $|T| \leq p-1$ and $|b| \leq p-1$. Let $X = (x_i)_{i\in[1..n]} \in \mathbb{Z}^n$ be the solution of $T.X = b$ over the integers. Then, $\forall\ k \in [0..n-1]$:*

$$\begin{cases} (p-2)^k - p^k \leq 2\frac{x_{n-k}}{p-1} \leq p^k + (p-2)^k & \text{if } k \text{ is even} \\ -p^k - (p-2)^k \leq 2\frac{x_{n-k}}{p-1} \leq p^k - (p-2)^k & \text{if } k \text{ is odd} \end{cases}$$

The proof is presented in [9, appendix A]. The idea is to use an induction on k with the relation $x_k = b_k - \sum_{i=k+1}^{n} T_{k,i}x_i$. Two lower and an upper bounds for x_{n-k} are computed, depending whether k is even or odd.

Corollary 3.3. $|X| \leq \frac{p-1}{2}\left[p^{n-1} + (p-2)^{n-1}\right]$.
Moreover, this bound is optimal.

Proof. We denote by $u_n = \frac{p-1}{2}\left[p^n - (p-2)^n\right]$ and $v_n = \frac{p-1}{2}\left[p^n + (p-2)^n\right]$ the bounds of the theorem 3.2. Now $\forall\ k \in [0..n-1]\ u_k \leq v_k \leq v_{n-1}$. Therefore the theorem 3.2 gives $\forall\ k \in [1..n]\ x_k \leq v_{n-1} \leq \frac{p-1}{2}\left[p^{n-1} + (p-2)^{n-1}\right]$

$$\text{Let } T = \begin{bmatrix} \ddots & \ddots & \ddots & \ddots & \ddots \\ & 1 & p-1 & 0 & p-1 \\ & & 1 & p-1 & 0 \\ & & & 1 & p-1 \\ & & & & 1 \end{bmatrix}, b = \begin{bmatrix} \vdots \\ 0 \\ p-1 \\ 0 \\ p-1 \end{bmatrix}$$

Then the solution $X = (x_i)_{i\in[1..n]} \in \mathbb{Z}^n$ of the system $T.X = b$ satisfies $\forall\ k \in [0..n-1]\ |x_{n-k}| = v_k$ □

Thus, for a given p, the dimension n of the system must satisfy

$$\frac{p-1}{2}\left[p^{n-1} + (p-2)^{n-1}\right] < 2^m \quad (1)$$

where m is the size of the mantissa so that the resolution over the integers using the BLAS trsm routine is exact. For instance, with a 53 bits mantissa, this gives quite small matrices, namely at most 55×55 for $p = 2$, at most 4×4 for $p \leq 9739$, and at most $p = 94906249$ for 2×2 matrices. Nevertheless, this technique is speed-worthy in most cases as shown in section 3.4.

3.3 Recursive with delayed modulus

In the previous section we noticed that BLAS routines within `Trsm` are used only for small systems. An alternative is to change the cascade: instead of calling the BLAS, one could switch to the classical iterative algorithm: Let $A \in {\mathbb{Z}_p}^{m\times m}$ and $B, X \in {\mathbb{Z}_p}^{m\times n}$ such that $AX = B$, then

$$\forall i, X_{i,*} = \frac{1}{A_{i,i}}(B_{i,*} - A_{i,[i+1..m]}X_{[i+1..m],*}) \quad (2)$$

The idea is that the iterative algorithm computes only one row of the whole solution at a time. Therefore its threshold t is greater than the one of the BLAS routine, namely it requires only

$$t(p-1)^2 < 2^m \quad (3)$$

Resultantly, an implementation of this iterative algorithm depends mainly on the matrix-vector product. The arithmetical cost of such an algorithm is now cubic in the size of

the system, where blocking improved the theoretical complexity. Anyway, in practice fast matrix multiplication algorithms are not better than the classical one for such small matrices [8, §3.3.2]. In section 3.4 we compare both hybrid implementations with different thresholds to the pure recursive one.

Now we focus on the dot product operation, base for matrix-vector product. We use the results of [6], extending those of [8, §3.1]. There several implementations of a dot product are proposed and compared on different architectures. According to [6], where many different implementations are compared (Zech log, Montgomery, float, ...), the best implementation is a combination of a conversion to floating point representation with delayed modulus (for big prime and vector size) and an overflow detection trick (for smaller prime and vector size).

The first idea is to specialize dot product in order to make several multiplications and additions before performing the division (which is then delayed). Indeed, one needs to perform a division only when the intermediate result might overflow. Now, if the available mantissa is of m bits and the modulo is p, divisions happen at worst every n multiplications where n satisfies condition (3). There the best compromise has to be chosen between speed of computation and available mantissa. A double floating point representation gives actually the best performances for most of the vector and prime sizes [6]. Moreover one can then perform the division "à la NTL" using a floating point precomputation of the inverse: $a * b \mod p = a * b - \lfloor a * b * p^{-1} \rfloor * p$.

For small primes, however, there is a faster method: the second idea is to use an integer representation and to let the overflow occur. Then one should detect this overflow and correct the result if needed. Indeed, suppose that we have added a product ab to the accumulated result t and that an overflow has occurred. The variable t now contains actually $t - 2^m$. Well, the idea is just to precompute a correction $CORR = 2^m \mod p$ and add this correction whenever an unsigned overflow has occurred. Now for a portable unsigned overflow detection, we use a trick of B. Hovinen[5]: since $0 < ab < 2^m$, an unsigned overflow has occurred if and only if $t + ab < t$. Of course, better performances are attained when several products are grouped (whenever possible) so that test and correction are also delayed [6]. Figure 1, shows the "quasi-optimal performances" obtained using both techniques on a Pentium 3 (P3): the constant curve is the very good behavior of a simple delayed division with a double floating point representation and "à la NTL" divisions. The step curve is a blocked version of the overflow and correction idea where a change of step reflects the need of an additional division. An optimal implementation would switch to one or the other representation depending on the size of the prime and on the architecture.

3.4 "Trsm" implementations behavior

As shown in section 3.1 the block recursive algorithm `Trsm` is based on matrix multiplications. This allows us to use the fast matrix multiplication routine of the FFLAS package [8]. This is an exact wrapping of the ATLAS library[6] used as a

[5]Personal communication, 2002

[6]`http://math-atlas.sourceforge.net`[22]

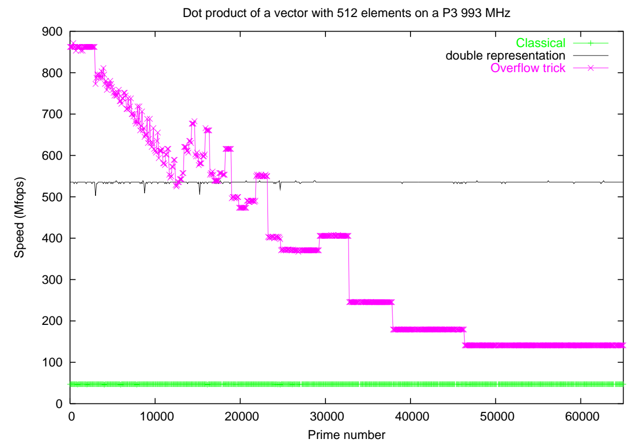

Figure 1: Speed improvement of dot product by delayed division, on a P3, 993 MHz

kernel to implement the `Trsm` variants. In the following we denote by "pure rec" the implementation of the recursive `Trsm` described in section 3.1. "BLAS" denotes the variant of section 3.2 with optimal threshold. "delayed$_t$" denotes the variant of section 3.3 where t satisfies equation 3. In our comparisons we use the fields presented in section 2 as base fields and the version 3.4.2 of ATLAS. Performances are expressed in million of finite field operations (`Mfops`) per second for $n \times n$ dense systems.

n	*400*	*700*	*1000*	*2000*	*3000*	*5000*
pure rec.	853	1216	1470	1891	2059	2184
BLAS	**1306**	**1715**	**1851**	**2312**	**2549**	**2660**
delayed$_{100}$	1163	1417	1538	1869	2042	2137
delayed$_{50}$	1163	1491	1639	1955	2067	2171
delayed$_{23}$	1015	1465	1612	2010	2158	2186
delayed$_3$	901	1261	1470	1937	2134	2166

n	*400*	*700*	*1000*	*2000*	*3000*	*5000*
pure rec.	810	1225	1449	1886	2037	2184
BLAS	1066	1504	**1639**	**2099**	**2321**	**2378**
delayed$_{100}$	1142	1383	1538	1860	2019	2143
delayed$_{50}$	**1163**	**1517**	**1639**	1955	2080	2172
delayed$_{23}$	1015	1478	1612	2020	2146	2184
delayed$_3$	914	1279	1449	1941	2139	2159

Table 1: Comparing speed (Mfops) of `Trsm` using `Modular<double>`, on a P4, 2.4GHz (Upper table is over Z_5, lower table is over Z_{32749})

One can see from table 1 that the "BLAS" `Trsm` variant with a `Modular<double>` representation is the most efficient choice for small primes (here switching to BLAS happens for $n = 23$ when $p = 5$ and $m = 53$). Now for big primes, despite the very small granularity (switching to BLAS happens only for $n = 3$ when $p = 32749$ and $m = 53$), this choice remains the best as soon as the systems are bigger than 1000×1000. This is because grouping operations into blocks speeds up the computation. Now in the case of smaller systems, the "delayed" variant is more efficient, due to the good behavior of dot product. However the threshold t has to be chosen carefully. Indeed using a threshold of 50 enables better performances than the BLAS variant. This is because the conversion from machine integers to floating point numbers becomes too big a price to pay. As a comparison, we

n	*400*	*700*	*1000*	*2000*	*3000*	*5000*
pure rec.	571	853	999	1500	1708	1960
BLAS	688	1039	1190	**1684**	**1956**	**2245**
delayed$_{150}$	799	**1113**	909	1253	1658	2052
delayed$_{100}$	**831**	1092	**1265**	1571	1669	2046
delayed$_{23}$	646	991	1162	1584	1796	2086
delayed$_{3}$	528	755	917	1369	1639	1903

n	*400*	*700*	*1000*	*2000*	*3000*	*5000*
pure rec.	551	786	1010	1454	1694	1929
BLAS	547	828	990	1449	1731	1984
delayed$_{100}$	703	958	1162	1506	1570	1978
delayed$_{50}$	**842**	**1113**	**1282**	**1731**	**1890**	**2174**
delayed$_{23}$	653	952	1086	1556	1800	2054
delayed$_{3}$	528	769	900	1367	1664	1911

Table 2: Comparing speed (Mfops) of `Trsm` using `Givaro-ZpZ`, on a P4, 2.4GHz (Upper table is over Z_5, lower table is over Z_{32749})

also provide performances for several thresholds, in particular the same as in the BLAS variant (3 and 23). Then for larger matrices, conversions ($O(n^2)$) are dominated by computations ($O(n^\omega)$), and the "BLAS" variant is again the fastest one, provided that the field is small enough.

Now, table 2 reflects experiments using the integer based modular fields of Givaro, `Givaro-ZpZ`. First, using these fields, `Trsm` is slower than with `Modular<double>`. This is due to the conversions needed for the matrix multiplications. Note that with the shown thresholds (not bigger than 100) the "overflow trick" used for the dotproduct slightly reduces this loss because it also reduces the number of conversions needed for the dotproduct.

To summarize, one would rather use a `Modular<double>` representation and the "BLAS" `Trsm` variant in most cases. However, when the base field is already specified, delayed$_t$ could provide slightly better performances. This requires a search for optimal thresholds which could be done through an Automated Empirical Optimizations of Software[22].

4. TRIANGULARIZATIONS

We now come to the core of this paper, namely the matrix multiplication based algorithms for triangularization over finite fields. The main concern here is the singularity of the matrices. Moreover, practical implementations need to efficiently deal with the rank profile, unbalanced dimensions, memory management, recursive thresholds, etc. Therefore, in this section we present three variants of the recursive exact triangularization. First the classical `LSP` of Ibarra et al. is sketched. In order to reduce its memory requirements, a first version, `LUdivine`, stores L in-place, but temporarily uses some extra memory. Our last implementation is fully in-place without any extra memory requirements and corresponds to Ibarra's `LQUP`. From both `LUdivine` and `LQUP` one can easily recover the `LSP` via some extractions and permutations.

4.1 LSP Factorization

The `LSP` factorization is a generalization of the well known block `LUP` factorization for the singular case [1]. Let A be a $m \times n$ matrix, we want to compute the triple $< L, S, P >$ such that $A = LSP$. The matrices `L` and `P` are as in `LUP` factorization and `S` reduces to a non-singular upper triangular matrix when zero rows are deleted. The algorithm with best known complexity computing this factorization uses a divide and conquer approach and reduces to matrix multiplication [16]. Let us describe briefly the behavior of this algorithm. The algorithm is recursive: first, it splits A in

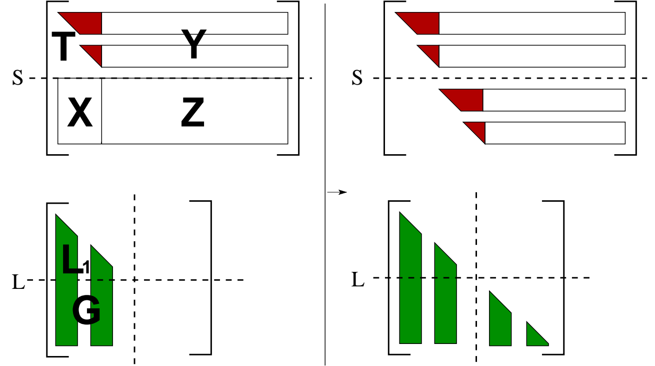

Figure 2: Principle of the LSP factorization

halves and performs a recursive call on the top block. It thus gives the T, Y and L_1 blocks of figure 2. Then, after some permutations ($[XZ] = [A_{21}A_{22}]P$), it computes G such that $GT = X$ via `Trsm`, replaces X by zeroes and eventually updates $Z = Z - GY$. The third step is a recursive call on Z. We let the readers refer e.g. to [2, (2.7c)] for further details.

Lemma 4.1. *Algorithm `LSP` is correct. The leading term of its theoretical cost is bounded by $\frac{C_\omega}{2^{\omega-1}-2} m^{\omega-1} \left(n + \frac{m}{2^{\omega}-2}\right)$ arithmetic operations in $\mathbb{Z}_p$ for $m \leq n$.*

This refines Ibarra's original factor[16, Theorem 2.1] from $3n$ to $n + \frac{m}{2^\omega-2}$. Moreover, when each one of the intermediate block is of full rank, this factor even reduces to $n - m\frac{2^{\omega-2}-1}{2^{\omega-1}-1}$[19, Theorem 1]. And this nicely gives $\frac{2}{3}n^3$, when $\omega = 3$, $n = m$ and $C_\omega = C_3 = 2$.

The point here is that, L being square $m \times m$ does not fit in place under S. Therefore a first implementation produces an extra triangular matrix. The following subsections address this memory issue.

4.2 LUdivine

The main concern with the direct implementation of the `LSP` algorithm, is the storage of the matrix L: it can not be stored with its zero columns under S (as shown in figure 2). Actually, there is enough room under S to store all the non zero entries of L, as shown in figure 3. Storing only the non zero columns of L is the goal of the `LUdivine` variant. One can notice that this operation corresponds to the storage of $\tilde{L} = LQ$ instead of L, where Q is a permutation matrix such that $Q^T S$ is upper triangular. Consequently, the recovery of L from the computed $\tilde{L}$ is straightforward. Note that this $\tilde{L}$ corresponds to the echelon form of [17, §2] up to some transpositions.

Further developments on this implementation are given in [3, 19]. However, this implementation is still not fully in place. Indeed, to solve the triangular system $G = X.T^{-1}$, one has then to convert T to an upper triangular matrix stored in a temporary memory space. In the same way, the matrix product $Z = Z - GY$ also requires a temporary

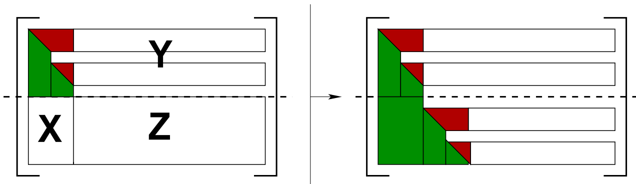

Figure 3: Principle of the LUdivine factorization

memory allocation, since rows of Y have to be shifted. This motivates the introduction of the `LQUP` decomposition.

4.3 LQUP

To solve the data locality issues, due to zero rows inside S, one can prefer to compute the `LQUP` factorization, also introduced in [16]. It consists in a slight modification of the `LSP` factorization: S is replaced by U, the corresponding upper triangular matrix, after the permutation of the zero rows. The tranpose of this row permutation is stored in Q.

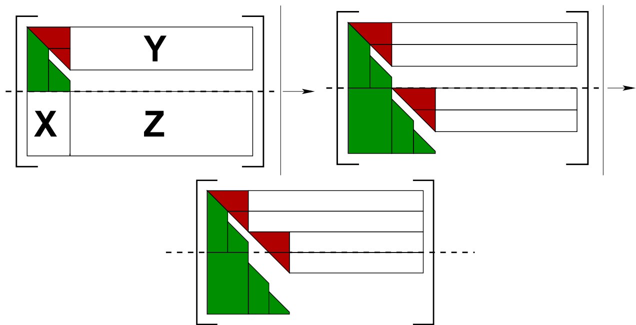

Figure 4: Principle of the LQUP factorization

This prevents the use of temporaries for Y and T, since the triangles in U are now contiguous. Moreover, the number of instructions to perform the row permutations is lower than the number of instructions to perform the block copies of `LUdivine` or `LSP`. Furthermore, our implementation of `LQUP` also uses the trick of `LUdivine`, namely storing L in its compressed form $\tilde{L}$. Thanks to all these improvements, this triangulation appears to be fully in place. As will be shown in section 4.4, it is also more efficient. Here again, the `LSP` and `LQUP` factorizations are simply connected via $S = QU$. So the recovery of the `LSP` is still straightforward.

4.4 Comparisons

As shown in previous sections the three variants of triangularization mainly differ by their memory management. Indeed, the main operations remain matrix multiplications and triangular system solving. Therefore, the implementation of all these variants use the fast matrix multiplication routine of the FFLAS package [8] and the triangular system solver of subsection 3.2 as kernel. The results are impressive: for example, table 3 shows that it is possible to triangularize a 5000×5000 matrix over a finite field in 29.9 seconds. We now compare the three routine speed and memory usage with the same kernels: a `Modular<double>` representation (so that no conversion overhead occur) and the recursive with BLAS `Trsm`. For table 3, we used random dense square

n	*400*	*1000*	*3000*	*5000*	*8000*	*10000*
`LSP`	0.05	0.48	8.01	32.54	404.8	1804
`LUdivine`	0.05	0.47	7.79	30.27	403.9	1691
`LQUP`	0.05	0.45	7.59	29.90	201.7	1090

Table 3: Comparing real time (seconds) of `LSP`, `LUdivine`, `LQUP` over $\mathbb{Z}_{101}$, on a P4, 2.4GHz

matrices (but with $3n$ non-zero entries) so as to have rank defficient matrices. The timings given in table 3 are close since the dominating operations of the three routines are similar. `LSP` is slower, since it performs some useless zero matrix multiplications when computing $Z = Z - GY$ (section 4.2). `LQUP` is slightly faster than `LUdivine` since row permutations involve less operations than the whole block copy of `LUdivine` (section 4.3). However these operations do not dominate the cost of the factorization, and they are therefore of little influence on the total timings. This is true until the matrix size induces some swapping, around 8000×8000.

Now for the memory usage, the fully in-place implementation of `LQUP` saves 20% of memory (table 4) when compared to `LUdivine` and 55% when compared to `LSP`. Actually, the memory usage of the original `LSP` is approximately that of `LUdivine` augmented by the extra matrix storage (which corresponds exactly to that of `LQUP`: e.g. $5000 * 5000 * 8bytes = 200Mb$). This memory reduction

n	*400*	*1000*	*3000*	*5000*	*8000*	*10000*
`LSP`	2.83	17.85	160.4	444.2	1136	1779
`LUdivine`	1.60	10.00	89.98	249.9	639.6	999.5
`LQUP`	1.28	8.01	72.02	200.0	512.1	800.0

Table 4: Comparing memory usage (Mega bytes) of `LSP`, `LUdivine`, `LQUP` over $\mathbb{Z}_{101}$, on a P4, 2.4GHz with 512 Mb RAM

is of high interest when dealing with large matrices (further improvements on the memory management are presented section 4.5).

4.5 Data locality

To solve even bigger problems, say that the matrices do not fit in RAM, one has mainly two solutions: either perform out of core computations or parallelize the resolution. In both cases, the memory requirements of the algorithms to be used will become the main concern. This is because the memory accesses (either on hard disk or remotely via a network) dominate the computational cost. A classical solution is then to improve data locality so as to reduce the volume of these remote accesses. In such critical situations, one may have to prefer a slower algorithm having a good memory management, rather than the fastest one, but suffering from high memory requirements. We here propose to deal with this concern in the case of rank or determinant computations of large dense matrices. The generalization to the full factorization case being direct but not yet fully implemented.

To improve data locality and reduce the swapping, the idea is to use square recursive blocked data formats [14]. A variation of the `LSP` algorithm, namely the `TURBO` algorithm

[10], adapts this idea to the exact case. Alike the LQUP algorithm which is based on a recursive splitting of the row dimension (see section 4.3), TURBO achieves more data locality by splitting both row *and* column dimensions. Indeed the recursive splitting with only the row dimension tend to produce "very rectangular" blocks: a large column dimension and a small row dimension. On the contrary, TURBO preserves the squareness of the original matrix for the first levels. More precisely each recursive level consists in a splitting of the matrix into four static blocks followed by five recursive calls to matrix triangularizations (U, V, C, D, and Z, in that order on figure 5), six Trsm and four matrix multiplications for the block updates. In this first implementation,

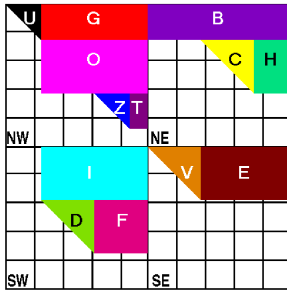

Figure 5: Principle of the TURBO decomposition

only one recursive step of TURBO is used, the five recursive calls being performed by the LQUP algorithm. For the actual size of matrices, the quite complex implementation of more recursive levels of TURBO is not yet mandatory.

Now for the comparisons of figure 6, we use the full LQUP factorization algorithm as a reference. Factorization of matrices of size below 8000 fit in 512Mb of RAM. Then LQUP is slightly faster than TURBO, implementation of the latter producing slightly more scattered groups. Now, the first field representation chosen (curves 1 and 2) is a modular prime field representation using machine integers. As presented in [8], any matrix multiplication occurring in the decomposition over such a representation is performed by converting the three operands to three extra floating point matrices. This memory overhead is critical in our comparison. TURBO, having a better data locality and using square blocks whenever possible, requires smaller temporary matrices than the large and very rectangular blocks used in LQUP. Therefore, for matrices of order over 8000, LQUP has to swap a lot while TURBO remains more in RAM. This is strikingly true for matrices between 8000 and 8500, where TURBO manages to keep its top speed.

Moreover, one can also reduce the memory overhead due to the conversions to floating point numbers, by using the modular<double> field representation, as described in section 2. There absolutely no allocation is done beside the initial matrix storage. On the one hand, performances increase since the conversions and copy are no longer performed, as long as the computations remain in RAM (see curves 3 and 4). On the other hand, the memory complexities of both algorithms now become identical. Furthermore, this fully inplace implementation does not create small block copies anymore. Paradoxically, this prevents the virtual blocks from

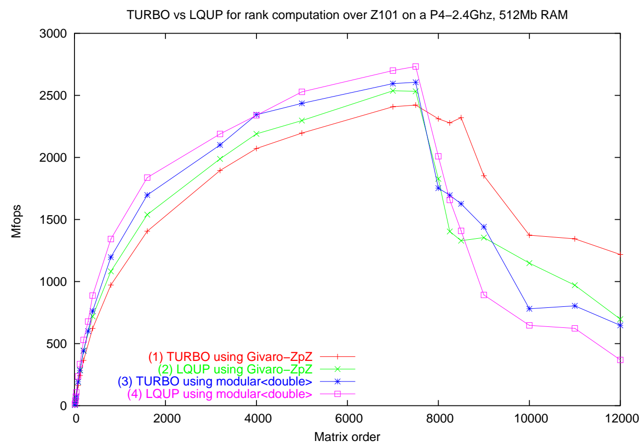

Figure 6: TURBO versus LQUP for out of core rank

fitting in the RAM, since they are just a view of the large initial matrix. For this reason, both performance losses appear for matrices of order around 8000. However, the drop is lower for TURBO thanks to the recursive blocked data formats producing better data locality.

This behavior of course confirms that as soon as the RAM is full, data locality becomes more important than memory saves : TURBO over Givaro-Zpz is the fastest for matrices of size bigger than 8000, despite its bigger memory demand. This is advocating further uses of recursive blocked data formats and of more recursive levels of TURBO.

5. RANK, DETERMINANT, INVERSE

The LQUP factorization and the Trsm routines reduce to matrix multiplication as we have seen in the previous sections. Theoretically, as matrix multiplication requires $2n^3 - n^2$ arithmetic operations, the factorization, requiring at most $\frac{2}{3}n^3$ arithmetic operations, could be computed in about $\frac{1}{3}$ of the time. Now, the matrix multiplication routine Fgemm of FFLAS package can compute 5000×5000 matrix multiplications in only 56.43 seconds on a 2.4GHz pentium 4. This is achieved with pipelining within the P4 processor and with very good performances of the BLAS. This corresponds to 4430 millions of finite field arithmetic operations per seconds! Well, table 5 shows that with $n \times n$ matrices we are not very far from these quasi-optimal performances also for the factorization:

Moreover, from the two routines, one can also easily derive

n	*400*	*700*	*1000*	*2000*	*3000*	*5000*
LQUP	0.05s	0.19s	0.45s	2.58s	7.59s	29.9s
Fgemm	0.05s	0.24s	0.66s	4.49s	12.66s	56.43s
Ratio	1	0.78	0.68	0.57	0.6	0.53

Table 5: Comparing cpu time of Matrix Multiplication and Factorization over $\mathbb{Z}_{101}$, on a P4, 2.4GHz

several other algorithms:

- The **rank** is the number of non-zero rows in U.
- The **determinant** is the product of the diagonal elements of U (stopping whenever a zero is encountered).
- The **inverse** is also straightforward:

Algorithm `Inverse`(A)
Input: $A \in \mathbb{Z}_p{}^{m \times m}$, non singular.
Output: $A^{-1} \in \mathbb{Z}_p{}^{m \times m}$.
Scheme

$$\begin{array}{lcl} L, U, P & := & \texttt{LQUP}(A). \quad (A \text{ is invertible, so } Q = I_m) \\ X & := & \texttt{LLeft-Trsm}(L, Id). \\ A^{-1} & := & P^T \, \texttt{ULeft-Trsm}(U, X). \end{array}$$

Now, the inverse can then be computed with at most $\frac{2}{3}n^3 + 2n^3$ arithmetic operations which gives a theoretical ratio of $\frac{4}{3}$. Once again, table 6 proves that our implementation has pretty good performances: Indeed, operations performed in

n	*400*	*700*	*1000*	*2000*	*3000*	*5000*
Inv	0.2s	0.76s	1.9s	11.27s	32.08s	132.72s
`Fgemm`	0.05s	0.24s	0.66s	4.49s	12.66s	56.43s
Ratio	4	3.17	2.92	2.51	2.53	2.35

Table 6: Comparing cpu time of Matrix Multiplication and Inverse over $\mathbb{Z}_{101}$, on a P4, 2.4GHz

`LQUP`, or `Trsm` are not grouped as well as in `Fgemm`. Therefore, the excellent performances of `Fgemm` make the ratio somewhat unreachable, although the `invert` routine is very fast. Note that, as the first `LLeft-Trsm` call is made on the identity, it could be accelerated in a specific routine. Indeed, during the course of LUP, L^{-1} can actually be computed with only a $\frac{n^3}{3}$ overhead, thus reducing the theoretical ratio from 4/3 to 1.

6. CONCLUSIONS

We have achieved the goal of approaching the speed of the numerical factorization of any shape and any rank matrices, but for finite fields. For example, the `LQUP` factorization of a 3000×3000 matrix over a finite field takes 7.59 seconds where 6 seconds are needed for the numerical `LUP` factorization of `lapack`[7]. To reach these performances one could use blocks that fit the cache dimensions of a specific machine. In [8] we proved that this was not mandatory for matrix multiplication. We think we prove here that this is not mandatory for any dense linear algebra routine. By the use of recursive block algorithms and efficient numerical BLAS, one can approach the numerical performances. Moreover, long range efficiency and portability are warranted as opposed to every day tuning with at most 10% loss for large matrices (see table 2 where delayed can beat BLAS only for big primes and with a specific empirical threshold).
Besides, the exact equivalent of stability constraints for numerical computations is coefficient growth. Therefore, whenever possible, we computed and improved theoretical bounds on this growth (see bounds 3.3 and [8, Theorem 3.1]). Those optimal bounds enable better uses of the BLAS.
Further developments include:
- A Self-adapting Software [5] (to switch to different algorithms during the recursive course of `Trsm` and `TURBO`), could be used to find the best empirical thresholds.
- The other case where our wrapping of BLAS is insufficient is for very small matrices (see tables 1 and 2). Here also, automated tuning would produce improved versions.
- The extension of the factorization to some other algorithms as shown for the Inverse (e.g. null-space computation as in [4]) is in progress.
- Finally, extending the out of core work of section 4.5 to design a parallel library is promising.

[7] `www.netlib.org/lapack`

Références

[1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *Design and Analysis of Computer Algorithms.* Addison-Wesley, 1974.

[2] D. Bini and V. Pan. *Polynomial and Matrix Computations, Vol. 1: Fundamental Algorithms.* Birkhauser, 1994.

[3] M. Brassel, P. Giorgi, and C. Pernet. LUdivine: A symbolic block LU factorization for matrices over finite fields using blas. ECCAD'2003, South Carolina, USA.

[4] Z. Chen and A. Storjohann. Effective reductions to matrix multiplication. ACA'2003, NC State University, USA.

[5] J. Dongarra and V. Eijkhout. Self-adapting numerical software and automatic tuning of heuristics. *Lecture Notes in Computer Science*, 2660:759–770, Jan. 2003.

[6] J.-G. Dumas. Efficient dot product over word-size finite fields. Rapport de recherche, IMAG-RR1064, Mar. 2004.

[7] J.-G. Dumas, T. Gautier, M. Giesbrecht, P. Giorgi, B. Hovinen, E. Kaltofen, B. D. Saunders, W. J. Turner, and G. Villard. LinBox: A generic library for exact linear algebra. ICMS'2002, Beijing, China, pp 40–50.

[8] J.-G. Dumas, T. Gautier, and C. Pernet. Finite field linear algebra subroutines. ISSAC'2002, Lille, France, pp 63–74.

[9] J.-G. Dumas, P. Giorgi, and C. Pernet. FFPACK: Finite Field Linear Algebra Package, preliminary version. Research Report, LIP-RR2004-02, Jan. 2004.

[10] J.-G. Dumas and J.-L. Roch. On parallel block algorithms for exact triangularizations. *Parallel Computing*, 28(11):1531–1548, Nov. 2002.

[11] J.-G. Dumas and G. Villard. Computing the rank of sparse matrices over finite fields. CASC'2002, Yalta, Ukraine, pp 47–62.

[12] P. Giorgi. From blas routines to finite field exact linear algebra solutions, July 2003. ACA'2003, NC State University, USA.

[13] P. Giorgi, C.-P. Jeannerod, and G. Villard. On the complexity of polynomial matrix computations. ISSAC'2003, Philadelphia, USA, pp 135–142.

[14] F. Gustavson, A. Henriksson, I. Jonsson, and B. Kaagstroem. Recursive blocked data formats and BLAS's for dense linear algebra algorithms. *Lecture Notes in Computer Science*, 1541:195–206, 1998.

[15] X. Huang and V. Y. Pan. Fast rectangular matrix multiplications and improving parallel matrix computations. PASCO'97, Maui, USA, pp 11–23.

[16] O. H. Ibarra, S. Moran, and R. Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms*, 3(1):45–56, Mar. 1982.

[17] E. Kaltofen, M. S. Krishnamoorthy, and B. D. Saunders. Parallel algorithms for matrix normal forms. *Linear Algebra and its Applications*, 136:189–208, 1990.

[18] C. Pernet. Implementation of Winograd's matrix multiplication over finite fields using ATLAS level 3 BLAS. Technical report, Laboratoire Informatique et Distribution, July 2001. `www-id.imag.fr/Apache/RR/RR011122FFLAS.ps.gz`.

[19] C. Pernet. Calcul du polynôme caractéristique sur des corps finis. Master's thesis, University of Delaware, 2003.

[20] C. Pernet and Z. Wan. LU based algorithms for the characteristic polynomial over a finite field. ISSAC'2003, Philadelphia, USA, pp 135–142. Poster.

[21] W. J. Turner. *Blackbox linear algebra with the LinBox library.* PhD thesis, NC State University, May 2002.

[22] R. C. Whaley, A. Petitet, and J. J. Dongarra. Automated empirical optimizations of software and the ATLAS project. *Parallel Computing*, 27(1–2):3–35, Jan. 2001.